arXiv:1611.09203v1 [cs.CV] 28 Nov 2016

Computational Mapping of the Ground Reflectivity with Laser Scanners

Juan Castorena *

Abstract

In this investigation we focus on the problem of mapping the ground reflectivity with multiple laser scanners mounted on mobile robots/vehicles. The problem originates because regions of the ground become populated with a varying number of reflectivity measurements whose value depends on the observer and its corresponding perspective. Here, we propose a novel automatic, data-driven computational mapping framework specifically aimed at preserving edge sharpness in the map reconstruction process and that considers the sources of measurement variation. Our new formulation generates map-perspective gradients and applies sub-set selection fusion and de-noising operators to these through iterative algorithms that minimize an ℓ_1 sparse regularized least squares formulation. Reconstruction of the ground reflectivity is then carried out based on Poisson's formulation posed as an ℓ_2 term promoting consistency with the fused gradient of map-perspectives and a term that ensures equality constraints with reference measurement data. We demonstrate our new framework outperforms the capabilities of existing ones with experiments realized on Ford's fleet of autonomous vehicles. For example, we show we can achieve map enhancement (i.e., contrast enhancement), artifact removal, de-noising and map-stitching without requiring an additional reflectivity adjustment to calibrate sensors to the specific mounting and robot/vehicle motion.

1. Introduction

Robotic mapping is one of the key fundamental tasks within the context of autonomous mobile robots and in particular in autonomous vehicles. The interest on this area is based on results that show that map augmentation can facilitate mobile robot/vehicle navigation and perception [1, 2]. Mapping of the environment is typically carried out by a mobile robot/vehicle equipped with position sensors (e.g., GPS, IMU, odometry) and range finders (e.g., lasers, radars, sonar) and/or cameras that perceive the environment while in motion. For example, [3] generates maps by capturing local orthographic projections of the ground with multiple vision cameras, [4–7] encodes a map of landmarks captured with a camera and matches them with those captured on the fly for localization. Unfortunately, capturing the outdoor environment with cameras is still susceptible to ambient illumination (e.g., unpredictable lighting) problems. LIDAR on the other hand rely on active illumination, does not significantly suffer from ambient illumination while also provides direct 3-D measurements of the environment. For these reasons, it is considered among one of the most reliable sensing modalities for 3-D mapping [8, 9].

Current LIDAR configurations are capable of measuring both the range and reflectivity for each pulse emitted. Here, reflectivity implies a factory calibration against a target of known reflectance has been performed. Specific benefits that accompany this additional modality include that of increased perception capabilities, for example to distinguish road features [10–14] in urban environments. One of the problems associated with mapping using this additional modality is that its measurements depend on a reflectivity response. This itself is dependent upon the laser-detector pair (i.e., observer), angle of incidence and range (i.e., perspective) [15]. Thus, assuming a

*J. Castorena (email: jcastore@ford.com) is with Ford Motor Company, Dearborn, MI, USA

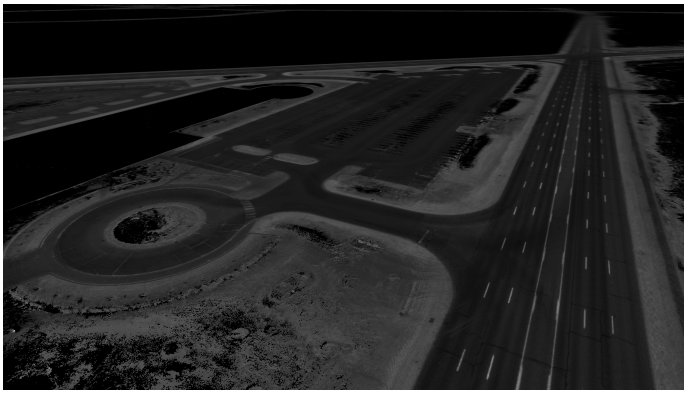

Figure 1. Map of the ground reflectivity (after full-SLAM).

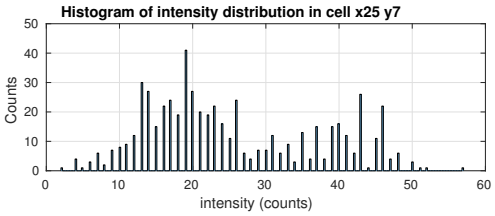

(a) Asphalt: $\mu \approx 25, \sigma \approx 12, D_{KL}(P||Q) = 0.38$ bits

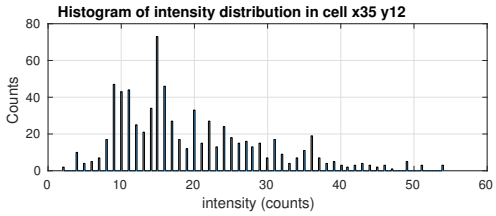

(b) W. lane: $\mu \approx 30, \sigma \approx 10, D_{KL}(P||Q) = 0.31$ bits

Figure 2. Histogram distributions of reflectivity measurements from multiple HDL-32E's in 10×10 cm. cells. (a) Cell in plain asphalt. (b) Cell in a white lane marking.

reflectivity invariant measurement model is quite often inadequate [16, 17] especially if one considers that the ground is viewed from multiple locations, perspectives and observers.

Typically, a map of the ground reflectivity is approximated as a uniformly quantized space plane of cell reflectivities. Each of these cells represent an area of size $\partial x \times \partial x$ (e.g., [9] uses $\partial x = 15$ cm). When a robot/vehicle is surveying a region, each reflectivity measurement of the ground is associated with a cell in the map depending on the projection of its corresponding range into the global reference frame (assuming full SLAM [18, 19] has already taken place). These projected measurements are non-uniformly distributed across the ground depending on the sensor configuration, the scanning pattern and the navigation path followed by the robot/vehicle while surveying. Naïve methods for estimating the cell reflectivity average all the raw reflectivity measurements associated to the cell, assuming these follow a Gaussian distribution model. However, as shown in Figure 2 and as measured by the Kullback-Leibler distance $D_{KL}(P||Q)$ this model is far from being true in practice.

Current mapping methods use an additional after-factory reflectivity calibration stage to decouple the observations from the variations in reflectivity [15]. Most of these approaches aim at estimating the response of each laser and its perspectives through a curve fit between the measurement signatures of a reference target and its corresponding known reflectance. For example, [20–23] exploit standard calibration targets and uses least squares to estimate the response of the laser scanner. Unfortunately, the options of target reflectances is very limited and measurements of these represent only a tiny fraction of the available data one may encounter in real world scenarios.

Within the context of autonomous vehicles, [9, 24] propose an automatic, target-less calibration method that uses a probabilistic framework to learn the reflectivity response model of each laser from its measurements. This method uses the expectation maximization iterative algorithm to generate the maximum likelihood model $p(y|x; b)$, that indicates the probability that for a given reflectivity x, laser b observes reflectivity y. However, to establish correspondences, the algorithm iterates between improving both the likelihood of the map $\mathbf{y}$ and the

laser responses. In practice, the whole process is time consuming and computationally intensive since it requires a data collection and its corresponding map generation step. Moreover, even after application of state of the art calibrations results in maps with artifacts, noise and of low contrast from global smoothing.

In this investigation we propose a novel computational mapping framework specifically aimed at preserving edge sharpness in the map reconstruction process. In particular, our approach first generates all the map-perspectives where each of these represent the map from a perspective of a laser observer. The problem with operating on these map-perspectives is that although these have overlapping observations of the same region, its perspective influences the characteristic reflectivity response. To decouple these from the measurement parameters we apply the gradient operator to each of them. Under this representation, the components of the gradients of map-perspectives become statistically stationary [25, 26] which allows one to operate on each invariantly. The algorithmic operators here proposed are those of perspective selection and de-noising which solve a corresponding ℓ_1 sparse regularized least squares cost function. Finally, the resulting map-perspective gradients are fused and used in a Poisson's formulation for map reconstruction of the ground reflectivity. The optimization carried out is posed as an ℓ_2 term promoting consistency with the gradient of fused map-perspective gradients and a term that ensures equality constraints with reference measured data.

2. Proposed Approach

2.1. Problem Formulation

The formulation we propose considers a map of the ground reflectivity compactly represented by the reflectivities $\mathbf{x} \in \mathbb{R}^{N_x \times N_y} \cup \{\infty\}^{N_x \times N_y}$ where N_x and N_y are the number of horizontal and vertical grid cells respectively. Here, the finite values in $\mathbf{x}$ represent the occupancy of measurements within the gridded map comprised of $N = N_x N_y$ cells. Note that in the remainder of this paper we often use $n \in \{1, \ldots, N\}$ to index the cells of $\mathbf{x}$ and other similarly sized matrices, essentially vectorizing them.

The laser reflectivity measurement model we use is the white Gaussian additive noise (AWGN) model

$$y^{\boldsymbol{\phi}} = \underbrace{f^b(x, \boldsymbol{\phi})}_{x^{\boldsymbol{\phi}}} + \epsilon^{\boldsymbol{\phi}} \tag{1}$$

where $f^b : \mathbb{R} \to \mathbb{R}$ is the non-linear response of laser with index b at $\boldsymbol{\phi}$ to a surface of reflectivity $x \in \mathbb{R}$, $x^{\boldsymbol{\phi}} \in \mathbb{R}$ is the surface reflectivity from the perspective of $\boldsymbol{\phi}$ and $\epsilon^{\boldsymbol{\phi}} \sim \mathcal{N}(0, \sigma_{\boldsymbol{\phi}})$ is an i.i.d. random variable independent of f^b. The set $\boldsymbol{\phi}$ indicates the laser observer and its perspective parameters; for example, $\boldsymbol{\phi} = (b, \theta, r)$ indicates the laser with global index b at an angle of incidence θ and range r. The set and number of globally projected (using full SLAM [18, 19] for example) reflectivity measurements of the ground at cell n taken from the perspective of $\boldsymbol{\phi}$ are denoted with $\tilde{\mathbf{y}}_n^{\boldsymbol{\phi}} \in \mathbb{R}^{M_n^{\boldsymbol{\phi}}}$ and $M_n^{\boldsymbol{\phi}}$, respectively. At a fixed $\boldsymbol{\phi}$, these measurements follow the AWGN model in (1) with the proper extension of terms, for example, the noise term follows $\boldsymbol{\epsilon}_n^{\boldsymbol{\phi}} \sim \mathcal{N}(\mathbf{0}_{M_n^{\boldsymbol{\phi}} \times M_n^{\boldsymbol{\phi}}}, \sigma_{\boldsymbol{\phi}} \cdot \mathbf{I}_{M_n^{\boldsymbol{\phi}} \times M_n^{\boldsymbol{\phi}}})$. From these measurements, we generate a set of map-perspectives $\mathbf{y}^{\boldsymbol{\phi}} \in \mathbb{R}^{N_x \times N_y} \cup \{\infty\}^{N_x \times N_y}$ which represent the map of ground reflectivity from the perspective of $\boldsymbol{\phi}$. These are defined element-wise as

$$[\mathbf{y}^{\boldsymbol{\phi}}]_n = \begin{cases} \frac{1}{M_n^{\boldsymbol{\phi}}} \sum_{m=1}^{M_n^{\boldsymbol{\phi}}} [\tilde{\mathbf{y}}_n^{\boldsymbol{\phi}}]_m, & \text{for} \quad n \in \Omega^{\boldsymbol{\phi}} \\ \infty & \text{for} \quad n \notin \Omega^{\boldsymbol{\phi}} \end{cases} \tag{2}$$

for all $\boldsymbol{\phi} \in \boldsymbol{\Phi}$ with $\boldsymbol{\Phi} = \{\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \ldots \boldsymbol{\phi}_B\}$ including as elements all of the possible perspective parameters from which measurements were collected and $B = |\boldsymbol{\Phi}| < \infty$ represents the total number of different perspectives from all observers. Examples of map-perspectives are shown in Figure 3 where the non-black regions represent the cell map occupancy $\Omega^{\boldsymbol{\phi}}$ of measurements from $\boldsymbol{\phi}$. The complete cell map occupancy or domain $\Omega' \subset \mathbb{R}^2$ is given

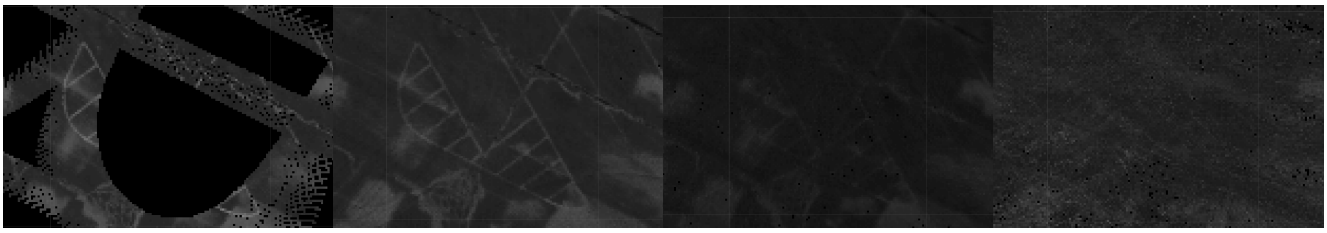

(a) Map-perspective of laser 1 (b) Map-perspective of laser 16 (c) Map-perspective of laser 19 (d) Map-perspective of laser 21

Figure 3. Ortographic view of patches of the ground reflectivity from the perspective of a given laser (indexed by order from bottom to top) in a Velodyne HDL-32E. Patches obtained while mapping with a vehicle in motion. (a) View from the perspective of laser 1, (b) View of beam 16: reduced contrast and full view of patch domain, (c) View of beam 19: showing bad contrast characteristic of high ranges and angles of incidence, (d) View of laser 21 presents the worst contrast.

by the union $\Omega' = \bigcup_{\boldsymbol{\phi}\in\boldsymbol{\Phi}} \Omega^{\boldsymbol{\phi}}$ of the map-perspective individual occupancies $\Omega^{\boldsymbol{\phi}}$. Note that by generating these map-perspectives a normalization to unbias for the non-uniform number of measurements from an observer and its perspectives in cells is carried out. Also, note that the uncertainty of the elements $[\mathbf{y}^{\boldsymbol{\phi}}]_n$ is now distributed as $\sim \mathcal{N}(0, \sigma_{\boldsymbol{\phi}}/M_n^{\boldsymbol{\phi}})$, essentially reducing noise at each cell by a factor of $M_n^{\boldsymbol{\phi}}$.

The problem of operating on these map-perspectives is that although these have overlapping observations of the same regions (i.e., overlapping occupancies), its perspective influences the characteristic reflectivity response. To decouple these from the perspective parameters we apply the gradient to each of them. Here, we denote the discrete gradient by $\mathbf{D} : \mathbb{R}^N \to \mathbb{R}^{N\times 2}$ composed of the matrices $\mathbf{D}_x : \mathbb{R}^N \to \mathbb{R}^N$ and $\mathbf{D}_y : \mathbb{R}^N \to \mathbb{R}^N$ that represent the first order finite forward difference along the horizontal and vertical axes, respectively. Under this representation, the gradients of map-perspectives components $[\mathbf{D}\mathbf{y}^{\boldsymbol{\phi}}]_n$ become statistically stationary which allows one to operate on each invariantly. The algorithmic operators here proposed are those of selection, de-noising and fusion which are compactly represented via the function

$$\widehat{\mathbf{D}\mathbf{y}} = f_{\text{fusion}}(\mathbf{y}^{\boldsymbol{\phi}_1}, \mathbf{y}^{\boldsymbol{\phi}_2}, ..., \mathbf{y}^{\boldsymbol{\phi}_B}). \tag{3}$$

where $\widehat{\mathbf{D}\mathbf{y}} \in \mathbb{R}^{N\times 2}$ denotes an estimate of the gradient of fused map-perspectives of $\mathbf{y}$. The goal of reconstruction on the other hand is to recover a map estimate $\widehat{\mathbf{x}} \in \mathbb{R}^{N_x \times N_y}$ of the ground reflectivity consistent with the fused gradients of map-perspectives and a reference measured data $\mathbf{x}^*$. Here, we compactly represent this reconstruction process by

$$\widehat{\mathbf{x}} = f_{\text{reconstruction}}(\mathbf{D}\mathbf{y}, \mathbf{x}^*). \tag{4}$$

Overall, our method is summarized in Figure 4.

In addition to this general formulation, we include here a few definitions. First, the sampling matrix $\mathbf{I}_\Omega \in \mathbb{R}^{|\Omega|\times N}$ is the submatrix of the identity obtained by extracting the rows indexed by Ω. This matrix is defined element-wise by

$$[\mathbf{I}_\Omega]_{i,j} = \begin{cases} 1 & i \in \{j \cap \Omega\} \\ 0 & \text{otherwise} \end{cases}. \tag{5}$$

Thus, a sampled vector $\mathbf{x}_\Omega \in \mathbb{R}^{|\Omega|}$ is the result of computing $\mathbf{x}_\Omega = \mathbf{I}_\Omega \mathbf{x}$. Second, the projection operator $\mathbf{P}_\Omega : \mathbb{R}^N \to \mathbb{R}^N$ projects inputs into the domain Ω. Application of this operator is performed element-wise as

$$[\mathbf{P}_\Omega \mathbf{x}]_n = \begin{cases} [\mathbf{x}]_n & n \in \Omega \\ 0 & \text{otherwise} \end{cases}. \tag{6}$$

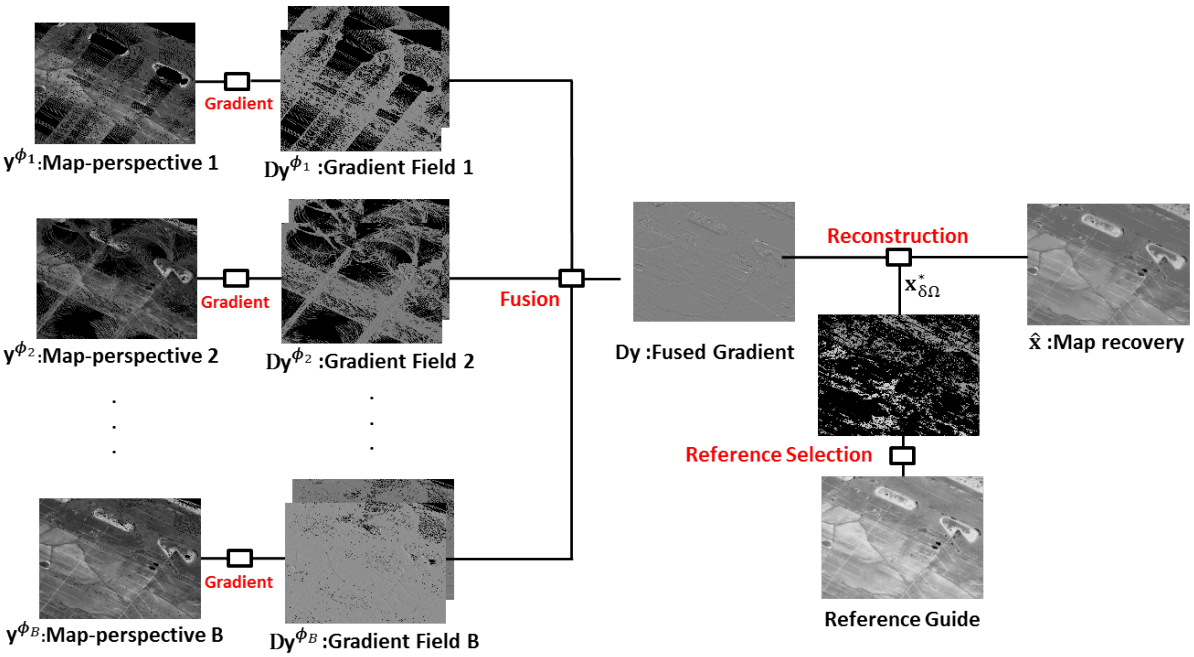

Figure 4. Schematic of our new method. The input images in the left column represent the map-perspectives obtained by using (2). The gradient field for each of these map-perspective images is then computed in the second column, followed by application of (3) and (4) within the fusion and reconstruction blocks, respectively.

Algorithm 1 Fusion of the gradient field with (8)

input: Map-perspective gradient $\mathbf{Dy}^{\boldsymbol{\phi}}$ for $\boldsymbol{\phi} \in \boldsymbol{\Phi}$, step $\gamma > 0$, and thresholding constant $\tau > 0$.
set: $t \leftarrow 1$, $q_0 \leftarrow 1$
repeat
  $\boldsymbol{s}^t \leftarrow \eta_\tau(\boldsymbol{s}^{t-1} - \gamma \nabla \mathcal{D}_1(\boldsymbol{s}^{t-1}))$
  $q_t \leftarrow \frac{1}{2}\left(1 + \sqrt{1 + 4q_{t-1}^2}\right)$
  $\mathbf{w}^t \leftarrow \boldsymbol{s}^t + ((q_{t-1} - 1)/q^t)(\boldsymbol{s}^t - \boldsymbol{s}^{t-1})$.
  $t \leftarrow t + 1$
until: stopping criterion
return: Fused gradient $\widehat{\mathbf{Dy}} = \sum_{\boldsymbol{\phi} \in \boldsymbol{\Phi}} \mathrm{w}^t_{\boldsymbol{\phi}} \cdot g_d(\mathbf{Dy}^{\boldsymbol{\phi}})$

Finally, for ease of notation, every element-wise operation that involves an ∞ value will be determined by ignoring it from its computation. For example,

$$\sum_{\boldsymbol{\phi} \in \boldsymbol{\Phi}} [\mathbf{y}^{\boldsymbol{\phi}}]_n = \sum_{\boldsymbol{\phi} \in \boldsymbol{\Phi}} [\mathbf{y}^{\boldsymbol{\phi}}]_n \cdot \mathbf{1}\{[\mathbf{y}^{\boldsymbol{\phi}}]_n < \infty\}. \tag{7}$$

2.2. Sparse selection, Denoising and Fusion

We propose here a fusion model of the map-perspectives in the gradient domain. This model consists in the weighted average of the map-perspective gradients $\mathbf{Dy}^{\boldsymbol{\phi}}$ obtained from each of the available $\mathbf{y}^{\boldsymbol{\phi}}$ (with ∞ values ignored as the example in (7)). In other words, fusion is carried out in the gradient domain by

$$\widehat{\mathbf{Dy}} = \sum_{\boldsymbol{\phi} \in \boldsymbol{\Phi}} \mathrm{w}_{\boldsymbol{\phi}} \cdot g_d(\mathbf{Dy}^{\boldsymbol{\phi}}). \tag{8}$$

where $g_d : \mathbb{R}^{N\times 2} \to \mathbb{R}^{N\times 2}$ can be any function including identity. The vector of weights $\mathbf{w} \in \mathbb{R}^{|\mathbf{\Phi}|}$ is chosen here to be a sparse vector that selects map-perspectives $\mathbf{y}^{\boldsymbol{\phi}}$ via its non-zero entries. In other words, we will find the best sparse representation of an overall averaged isotropic gradient of map-perspectives. The reason for this is that some map-perspectives should be penalized or even removed from the reconstruction process as illustrated in Figure 3. To achieve this, we find the weights that minimize the isotropic sparse promoting least squares optimization problem

$$\widehat{\mathbf{w}} = \arg\min_{\mathbf{w}\in\mathcal{W}} \left\{ \mathcal{D}_1(\mathbf{w}) + \lambda \mathcal{R}(\mathbf{w}) \right\}, \tag{9}$$

where minimization is carried out over a fidelity term $\mathcal{D}_1 : \mathbb{R}^{|\mathbf{\Phi}|} \to \mathbb{R}$ and a regularization term $\mathcal{R} : \mathbb{R}^{|\mathbf{\Phi}|} \to \mathbb{R}$ controlled by the parameter $\lambda > 0$ which determines the sparseness strength. In this problem the fidelity term is defined by

$$\mathcal{D}_1(\mathbf{w}) = \frac{1}{2} \left\| \sum_{\boldsymbol{\phi}\in\mathbf{\Phi}} |\mathbf{D}\mathbf{y}^{\boldsymbol{\phi}}| - \sum_{\boldsymbol{\phi}\in\mathbf{\Phi}} \mathrm{w}_{\boldsymbol{\phi}} |\mathbf{D}\mathbf{y}^{\boldsymbol{\phi}}| \right\|_{\ell_2}^2 \tag{10}$$

which promotes consistency between the weighted average and the average of map-perspective gradients. The second regularization term

$$\mathcal{R}(\mathbf{w}) = \|\mathbf{w}\|_{\ell_1} \tag{11}$$

promotes sparse solutions. Adjustment of the parameter λ can be useful to control the number of map-perspectives $\mathbf{y}^{\boldsymbol{\phi}}$ that will go into the fusion process and map reconstruction.

The problem in (9) involving the non-convex functional in (11) can be solved iteratively using a proximal gradient algorithm. This involves the alternate application of a gradient-descent step on the fidelity term in (10) followed by the application of the non-linear soft-thresholding operator $\eta_\tau(\cdot) : \mathbb{R}^N \to \mathbb{R}^N$ defined element-wise by

$$[\eta_\tau(\mathbf{x})]_n = \mathrm{sgn}([\mathbf{x}]_n)(|[\mathbf{x}]_n| - \tau)_+ \tag{12}$$

to promote sparsity of the solution at each iteration. The accelerated implementation of this algorithm is known as fast iterative shrinkage thresholding algorithm (FISTA) [27] which converges with a rate of $O(1/k^2)$. A summary of the selection algorithm is given in Algorithm 1.

In (8), we include the function g_d and propose the option to use it for denoising using the non-linear soft-thresholding operator [28] applied independently to each of the horizontal and vertical gradient components. The basis of this idea can be found in [29] and our proposed implementation is summarized in Algorithm 3 in the Appendix.

2.3. Map reconstruction

The next step consists of the reconstruction of the map of ground reflectivity based on the output of the fusion process. This step is required since we are interested in reconstructing a map of reflectivity and not of gradients. Specifically, we use $\widehat{\mathbf{D}\mathbf{y}}$ from (3) to compactly denote an estimate of the fused gradient of map-perspectives. Reconstruction is then carried out by minimizing a cost function on $\mathbf{x}$ based on Poisson's formulation [30]. Here, this problem is posed as an ℓ_2 term promoting consistency with the fused gradient of map-perspectives and a term that ensures equality constraints with reference measurement data on a subset $\partial\Omega \subset \Omega'$ of cells of $\mathbf{x}$. Based on this descriptions the map of ground reflectivity can be recovered by solving

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}\in\mathcal{X}} \left\{ \frac{1}{2} \sum_{k=\{x,y\}} \|\mathbf{D}_k\widehat{\mathbf{y}} - \mathbf{D}_k\mathbf{x}\|_{\ell_2}^2 \right\}, \text{ s.t. } \mathbf{x}_{\partial\Omega} = \mathbf{x}^*_{\partial\Omega} \tag{13}$$

Algorithm 2 Map reconstruction of ground reflectivity

input: Fused gradient $\mathbf{Dy}$, reference $\mathbf{x}^*_{\partial\Omega}$ and $\gamma > 0$.
set: $t \leftarrow 1$
repeat
  $\boldsymbol{s}^t \leftarrow \boldsymbol{s}^{t-1} - \gamma \nabla \mathcal{D}_2(\boldsymbol{s}^{t-1})$
  $q_t \leftarrow \frac{1}{2}\left(1 + \sqrt{1 + 4q_{t-1}^2}\right)$
  $\mathbf{x}^t_\Omega \leftarrow \boldsymbol{s}^t + ((q_{t-1} - 1)/q^t)(\boldsymbol{s}^t - \boldsymbol{s}^{t-1})$.
  $t \leftarrow t + 1$
until: stopping criterion
set: $\widehat{\mathbf{x}}_{\partial\Omega} \leftarrow \mathbf{x}^*_{\partial\Omega}, \widehat{\mathbf{x}}_\Omega \leftarrow \mathbf{x}^t_\Omega$
return: Map of ground reflectivity $\widehat{\mathbf{x}}$

The minimizer of (13) is the unique solution to the equivalent Poisson equation with Dirichlet boundary conditions problem posed in matrix form as:

$$\mathbf{I}_\Omega \mathbf{L}\mathbf{x} = \mathbf{I}_\Omega[\mathbf{D}_x^2\mathbf{y} + \mathbf{D}_y^2\mathbf{y}], \quad \text{s.t.} \quad \mathbf{x}_{\partial\Omega} = \mathbf{x}^*_{\partial\Omega} \tag{14}$$

where $\mathbf{L} : \mathbb{R}^N \to \mathbb{R}^N$ is a discrete Laplacian operator defined in (19) in the appendix. To simplify this problem, we recast (14) in a form such that optimization is carried out only on the entries in the domain of $\Omega = \partial\Omega^c$ representing the cell locations where reconstruction is to be performed, since we already know $\mathbf{x}_{\partial\Omega}$. To accomplish this, we use the projection and sampling operators defined in (6) and (5), respectively, and rewrite (14) as the minimization problem:

$$\widehat{\mathbf{x}}_\Omega = \underset{\mathbf{x}_\Omega \in \mathcal{X}_\Omega}{\arg\min} \left\{\mathcal{D}_2(\mathbf{x}_\Omega)\right\}, \quad \text{and} \quad \widehat{\mathbf{x}}_{\partial\Omega} = \mathbf{x}^*_{\partial\Omega} \tag{15}$$

where $\mathcal{D}_2 : \mathbb{R}^{|\Omega|} \to \mathbb{R}$ represents the fidelity term that satisfies (14) over the domain Ω. In other words,

$$\mathcal{D}_2(\mathbf{x}) = \frac{1}{2}\left\|\mathbf{\Phi}\mathbf{x} - \mathbf{b}\right\|_{\ell_2}^2 \tag{16}$$

with $\mathbf{\Phi}$ representing the Ω projected and sampled Laplacian of the left hand side of the first equality in (14) and $\mathbf{b}$ represents its right side with incorporation of the Laplacian projection of the second equality contraint. Mathematically, these two terms are defined via

$$\mathbf{\Phi} = \mathbf{I}_\Omega \mathbf{L}\mathbf{P}_\Omega \mathbf{I}_\Omega^T \quad \text{and} \quad \mathbf{b} = \mathbf{I}_\Omega[\mathbf{D}_x^2\mathbf{y} + \mathbf{D}_y^2\mathbf{y} - \mathbf{L}\mathbf{P}_{\partial\Omega}\mathbf{x}^*] \tag{17}$$

where the superscript ^{T} denotes a matrix transpose. Since the cost functional in (16) is convex, we rely on the accelerated gradient projection method of Nesterov [31] which achieves a rate of convergence of $O(1/k^2)$. This iterative method is summarized in Algorithm 2.

To define $\mathbf{x}^*_{\partial\Omega}$ and the set $\partial\Omega$ there are several alternatives that could be used. Here, the method we propose simply consists on choosing a map-perspective $\mathbf{y}^{\boldsymbol{\phi}_1}$ characterized with a strong gradient compared to others within some region (e.g., $|[\mathbf{D}(\mathbf{y}^{\boldsymbol{\phi}_1})]_n| > |[\mathbf{D}(\mathbf{y}^{\boldsymbol{\phi}})]_n|$; for some $n \in \Omega^{\boldsymbol{\phi}_1} \cap \Omega^{\boldsymbol{\phi}}; \forall \boldsymbol{\phi} \neq \boldsymbol{\phi}_1$ if it exists). From the map-perspective $\mathbf{y}^{\boldsymbol{\phi}_1}$, we extract the subset of cells indexed by $\partial\Omega$ and use the reflectivity values in the corresponding cells to construct our reference as $\mathbf{x}^*_{\partial\Omega} = \mathbf{y}^{\boldsymbol{\phi}_1}_{\partial\Omega}$. The set $\partial\Omega$ is determined from $\mathbf{y}^{\boldsymbol{\phi}_1}$ as the cells with the most-likely reflectivity value, in other words $\partial\Omega = \{n | [\mathbf{y}^{\boldsymbol{\phi}_1}]_n = \underset{y}{\arg\max}\, p_{\mathbf{y}^{\boldsymbol{\phi}_1}}(\mathbf{y}^{\boldsymbol{\phi}_1})\}$.

3. Experiments

To verify our approach, we use mapping datasets from [32] and a few additional we collected with Ford's Fusion testing fleet of autonomous vehicles. These are outfitted with four Velodyne HDL-32E 3D-LIDAR scanners and an

Applanix POS-LV 420 (IMU). The four HDL-32E's were mounted on the roof above the B pillars of the vehicle; two which rotate in an axis perpendicular to the ground (i.e., non-canted) and two that are canted. Extrinsic calibration of the LIDAR's is performed using GICP [33]. When mapping, all LIDAR measurements are projected into a global reference frame using the vehicle pose from the GPS, IMU and odometry with a post-correction by full-SLAM [34, 35]. The datasets we used were all collected at the Ford campus in Dearborn, Michigan. To test our approach we use only data from the two non-canted LIDARs. Under this configuration, the angle of incidence and range per laser recorded from the ground remains approximately constant over the course of the 360°scan trip. Thus, each laser index implicitly accounts for these two which then results in a total of $B = 64$ map-perspectives. To compare against state of the art, we test against an implementation of the deterministic reflectivity calibration method in [9,24]. The data we used in this case is collected from all four LIDARs. The two main differences in our implementation are: (1) we make an approximation to the generate the look-up table in one data pass and (2) we include an angle of incidence dependency for the canted LIDARs. We would like to mention that the calibration and map generation dataset is the same thus implying the best possible map results.

The implementation of our method is applied independently on disjoint patches of the map of size 400×400 cells with 10×10 cm cells. First, the search spaces are defined by $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^N : 0 \leq \mathbf{x} \leq 255\}$, $\mathcal{X}_\Omega = \{\mathbf{x}_\Omega \in \mathbb{R}^{|\mathbf{\Phi}|} : 0 \leq \mathbf{x}_\Omega \leq 255\}$, $\mathcal{X}_\mathbf{D} = \{\mathbf{D}_k\mathbf{x} \in \mathbb{R}^N : -255 \leq \mathbf{D}_k\mathbf{x} \leq 255\}$. Second, we find that Algorithm 1 converges in $K \leq 100$ iterations with $\lambda = 1.2e-3$, $\gamma = 1e^{-3}$ and $\tau = 2.3e^{-3}$. In the denoising Algorithm 3 we use the same parameter settings as in Algorithm 1. For the reconstruction in Algorithm 2, we found that the relative energy step in two successive iterations $\|x_\Omega^t - x_\Omega^{t-1}\|_{\ell_2}/\|x_\Omega^{t-1}\|_{\ell_2} < 1e^{-3}$ or $K \leq 256$ iterations is sufficient to yield good reconstructions. To give an intuition of the run-times, it took 45 secs to reconstruct a patch of size 40×40 meters.

3.1. Contrast enhancement

In this subsection we show that our method results in maps of enhanced contrast. Figure 5 includes a comparison of map patches between raw naïve fusion (i.e., average of the raw reflectivity measurements) in Figure 5.a, our implementation of state of the art calibration [9] in Figure 5.c and the result of our approach applied to the two corresponding cases in Figure 5.b and d. Note that our approach yields in general significantly better contrast (e.g., smoother regions with sharper edges). Even some of the edges from lane markings which were close to being indistinguishable appear more visible with our method.

3.2. Artifact removal

We have found that under certain driving conditions the non-uniform laser reflectivity response results in map patches corrupted with artifacts. These have been observed also in other works, for example even after applying the reflectance calibration in [24, 36]. In this section, we include examples in which we show three types of artifacts.

Figure 6.a points in red the circular artifacts that are generated close to the middle of the road. This artifact pattern coincides with the 360° laser scanning pattern while the vehicle is undergoing motion at longitudinal velocity > 6 m/s. Figure 6.b shows the generated result when reconstructed with the calibration method of [24]. This method eliminated most of the artifact, although, at the cost of global smoothing. Figure 6.c shows the result obtained with our method. Note that a great deal of the artifact is removed while also enhancing (with respect to 6.a and 6.b) edge sharpness.

Figure 6.d shows an artifact example generated when the vehicle stopped for more than > 5 seconds at an intersection. The artifact follows the pattern of the lasers while undergoing a 360° scan trip and is caused by cells accumulating measurements from lasers with dominant reflectivity responses. Figure 6.e shows the result of applying our implementation of [24]. Note, that the global smoothing removed a great deal of the artifact, although not in its entirety. Finally, Figure 6.f illustrates the map with sharp edges and removed artifact obtained with our reconstruction approach.

The last case shown in Figure 6.g presents an artifact generated when the vehicle underwent a u-turn. Figure

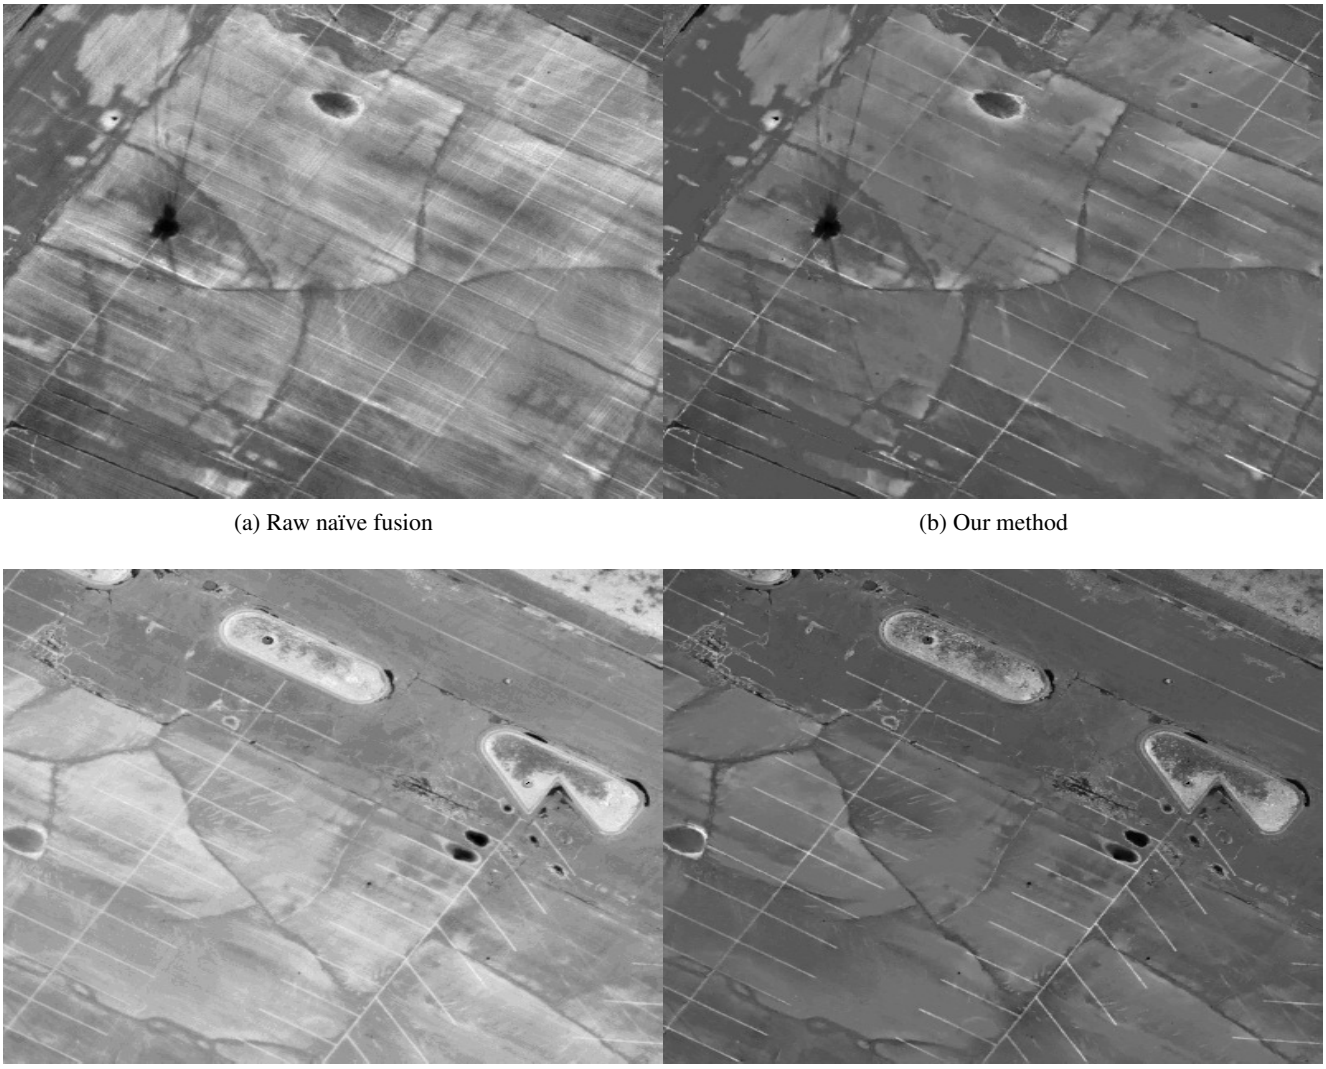

(a) Raw naïve fusion (b) Our method

(c) Reflectivity calibration based (d) Our method

Figure 5. Contrast enhancement in map patches with ice on the ground. Note: gray scale value represents reflectivity. **Zoom in to better appreciate image quality.**

6.h shows the result obtained with our practical implementation of [24]. The artifact generated in the u-turn is eliminated, but again at the cost of global smoothing. In contrast, our method illustrated in Figure 6.i presents smooth areas while maintaining sharp edges and eliminating the u-turn artifact.

3.3. Sparse selection and Denoising

In this subsection we include results that show the plausible operators that can be applied to render maps with improved quality characteristics at the cost of a computational complexity trade-off. For example, the possiblity of sparsely selecting the map-perspectives using (9) and/or include the denoising step through g_d in (8). Here, we include a representative example that illustrates the differences between the quality of maps in four cases: (1) Uniform fusion weights (i.e., $w_{\boldsymbol{\phi}} = 1$) and no denoising, (2) Uniform fusion weights with denoising, (3) Selection of 5 map-perspectives through control of $\lambda = 1.2e - 3$ in (9) with no denoising and (4) Same as (3) but with de-noising. Figure 7 shows a comparison between the aforementioned cases and the raw naïve fusion in Figure 7.a and our implementation of the calibration of [9] in Figure 7.b. The obtained results with the adjustments

(a) Raw naïve fusion (b) Implementation of Levinson & Thrun [24] (c) Our method

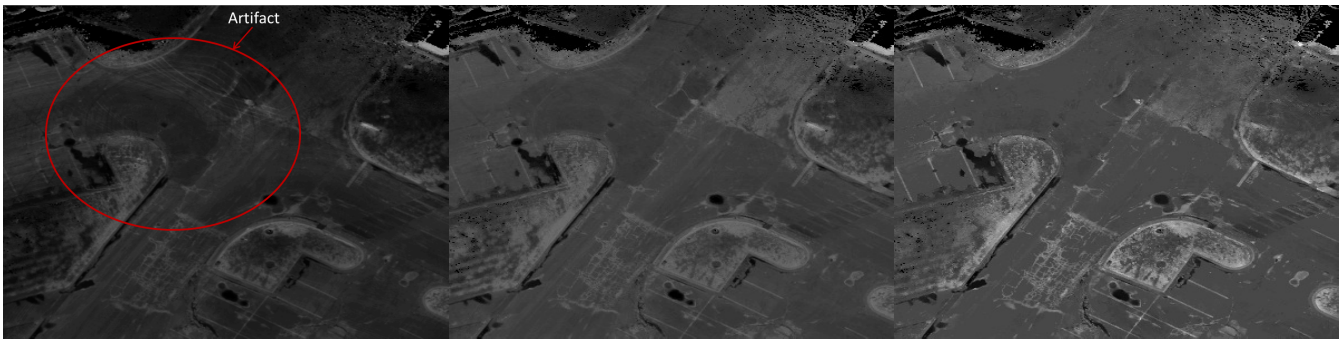

(d) Raw naïve fusion (e) Implementation of Levinson & Thrun [24] (f) Our method

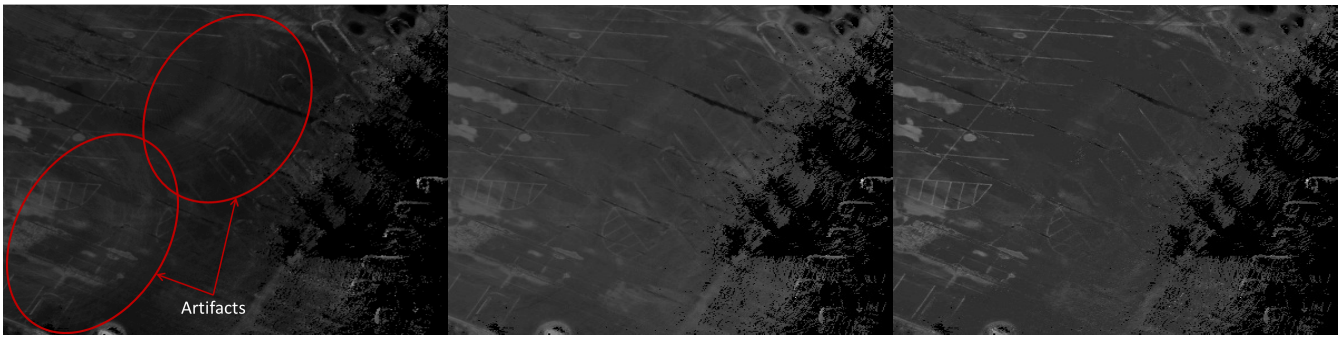

(g) Raw naïve fusion (h) Implementation of Levinson & Thrun [24] (i) Our method

Figure 6. Patch comparison of artifact correction. (a,d,g) Raw naïve fusion, note: red lines indicate location of artifacts. (b,e,h) Results of our implementation of the reflectivity calibration in [24], (c,f,i) Our method. **Note: Zoom in to better appreciate image quality.**

(1),(2), (3) and (4) are included in Figures 7.c-f, respectively. Note that even the result in Figure 7.c obtained when averaging all map-perspectives with no denoising presents smoother regions with sharper edges in comparison to Figure 7.a-b. Figure 7.d adds more smoothing while preserving edge sharpness to 7.c with the denoising step. Finally, Figures 7.e and 7.f illustrate that the sparse selection in the two cases: no denoising and with denoising, respectively results in maps with improved edge sharpness

4. Conclusion

In this investigation we proposed a novel computational mapping framework specifically designed to preserve edge sharpness in maps generated with data from robots/vehicles equipped with laser scanners. We showed

Algorithm 3 Anisotropic FISTA for denoising gradients.

input: Gradient $\mathbf{Dy}$, step γ, and thresholding const. τ.
for each component field $k \in \{x, y\}$
  set: $t \leftarrow 1$, $s^0 \leftarrow \mathbf{D}_k\mathbf{y}$
  repeat
    $\boldsymbol{s}^t \leftarrow \eta_\tau(\boldsymbol{s}^{t-1} - \gamma\nabla\mathcal{D}_3(\boldsymbol{s}^{t-1}))$
    $q_t \leftarrow \frac{1}{2}\left(1 + \sqrt{1 + 4q_{t-1}^2}\right)$
    $\mathbf{D}_k\mathbf{x}^t \leftarrow \boldsymbol{s}^t + ((q_{t-1} - 1)/q^t)(\boldsymbol{s}^t - \boldsymbol{s}^{t-1})$.
    $t \leftarrow t + 1$
  until: stopping criterion
return: Denoised gradient of map-perspective $\mathbf{Dx}^t$

through experimentation that our edge-preserving method enhances contrast, removes artifacts, denoises and has map-sticthing capabilities that outperform those of existing calibration based ones. The main reason of this being that our method optimizes the estimate of the map of ground reflectivity instead of the individual laser responses. Moreover, our gradient based fusion approach removes the requirement of any post factory reflectivity calibration. This represents a significant advantage if one considers the cost implied in calibration against standard reference targets and/or the time requirements of post-factory calibration processes, which are currently unfeasible for robots/vehicles produced in mass.

A. Appendix

Assuming a column-wise vectorization of a 2D signal of size $N_y \times N_x$, the first order forward difference discrete gradient is defined point-wise as

$$[\mathbf{Dx}]_n = \begin{pmatrix} [\mathbf{D}_x\mathbf{x}]_n \\ [\mathbf{D}_y\mathbf{x}]_n \end{pmatrix} = \begin{pmatrix} [\mathbf{x}]_{n+N_y} - [\mathbf{x}]_n \\ [\mathbf{x}]_{n+1} - [\mathbf{x}]_n \end{pmatrix}. \tag{18}$$

were $\mathbf{D}_x$ and $\mathbf{D}_y$ represent the horizontal and vertical components, respectively. In a similar way, the discrete Laplacian operator $\mathbf{L} : \mathbb{R}^N \to \mathbb{R}^N$ is defined point-wise as

$$[\mathbf{Lx}]_n = -4[\mathbf{x}]_n + [\mathbf{x}]_{n+1} + [\mathbf{x}]_{n-1} + [\mathbf{x}]_{n+N_y} + [\mathbf{x}]_{n-N_y} \tag{19}$$

A.1. Soft-tresholding algorithm

To denoise a map-perspective, we propose to apply the soft-thresholding [28] to its gradient-field. This method solves the sparse promoting least squares optimization

$$\mathbf{D}_k\mathbf{x} = \arg\min_{\mathbf{D}_k\mathbf{x}\in\mathcal{X}_\mathbf{D}} \left\{\mathcal{D}_3(\mathbf{D}_k\mathbf{x}) + \lambda\mathcal{R}(\mathbf{D}_k\mathbf{x})\right\}, \tag{20}$$

for each horizontal and vertical $k = \{x, y\}$ direction, independently. Here, $\lambda > 0$ controls the amount of regularization (i.e., sparsity). The first term in (20) measures the gradient fidelity defined in the least squares sense as

$$\mathcal{D}_3(\mathbf{D}_k\mathbf{x}) = \frac{1}{2}\left\|\mathbf{D}_k\mathbf{y} - \mathbf{D}_k\mathbf{x}\right\|_{\ell_2}^2 \tag{21}$$

while the second term is the non-smooth ℓ_1 sparse promoting regularizer defined by

$$\mathcal{R}(\mathbf{D}_k\mathbf{x}) = \|\mathbf{D}_k\mathbf{x}\|_{\ell_1}. \tag{22}$$

The optimization in (20) can be iteratively solved using the accelerated gradient descent in [31] along with the non-convex proximity projection method of [37]. The complete Algorithm in 3 has a rate of convergence of $O(1/k^2)$.

References

[1] Google, "Google self-driving car project monthly report," https://static.googleusercontent.com/selfdrivingcar/files/reports/report-0316.pdf, 2016, [Online: accessed 21-September-2016]. 1

[2] Popular Mechanics, "Whoever owns the maps owns the future of self-driving cars," http://www.popularmechanics.com/cars/a21609/here-maps-future-of-self-driving-cars/, 2016, [Online: accessed 21-September-2016]. 1

[3] Napier A. and Newman P., "Generation and exploitation of local orthographic imagery for road vehicle localisation," in *IEEE Intelligent Vehicles Symposium (IV)*, July 2012, vol. IV. 1

[4] S. Se, D. Lowe, and J. Little, "Vision-based mobile robot localization and mapping using scale-invariant features," in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001, pp. 2051–2058. 1

[5] S. Se, D. Lowe, and J. Little, "Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks," *The International Journal of Robotics Research*, vol. 21, no. 8, pp. 735–758, August 2002. 1

[6] S. Se, D. Lowe, and J. Little, "Vision-based global localization and mapping for mobile robots," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 364–375, June 2005. 1

[7] N. Karlsson, E. Di Bernardo, J. Ostrowski, L. Goncalves, P. Pirjanian, and M.E. Munich, "The vslam algorithm for robust localization and mapping," in *IEEE International Conference on Robotics and Automation*, 2005. 1

[8] J. Levinson, Montemerlo M., and S. Thrun, "Map-based precision vehicle localization in urban environments," in *Robotics Science and Systems*, 2007. 1

[9] J. Levinson and S. Thrun, "Robust vehicle localization in urban environments using probabilistic maps," *IEEE International Conference on Robotics and Automation*, pp. 4372–4378, May 2010. 1, 2, 8, 9

[10] H. Guan, J. Li, Y. Yu, C. Wang nad M. Chapman, and B. Yang, "Using mobile laser scanning data for automated extraction of road markings," *ISPRS Journal of Photogrammetry and Remote Sensing*, , no. 87, pp. 93–107, December 2015. 1

[11] M.J. Olsen, F. Kuester nad B.J. Chang, and T.C. Hutchinson, "Terrestrial laser scanning-based structural damage assessment," *Journal of Computing in Civil Engineering*, vol. 24, no. 3, 2010. 1

[12] H. Guan, Y. Yu, J. Li, P. Liu, H. Zhao, and C. Wang, "Automated extraction of manhole covers using mobile lidar data," *Remote sensing letters*, vol. 5, pp. 1042–1050, 2014. 1

[13] S. Pu, M. Rutzinger, G. Vosselman, and S. O. Elberink, "Recognizing basic structures from mobile laser scanning data for road inventory studies," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 66, no. 6, pp. S28–S39, 2011. 1

[14] A. Chin and M.J. Olsen, "Evaluation of technologies for road profile capture, analysis, and evaluation," *Journal of Surveying Engineering*, vol. 141, no. 1, 2014. 1

[15] A. Kashani, M. Olsen, C.E. Parrish, and N. Wilson, "A review of lidar radiometric processing: From ad hoc intensity correction to rigorous radiometric calibration," *Sensors*, , no. 15, pp. 28099–28128, November 2015. 1, 2

[16] B.R. Harvey and D.D. Lichti, "The effects of reflecting surface material properties on time-of-flight laser scanning measurements," *International properties on time-of-flight laser scanning measurements.*, 2005. 2

[17] W. Burgard, R. Triebel, and H. Andreasson, "Improving plane extraction from 3d data by fusing laser data and vision," *Intelligent Robotics and Systems*, 2005. 2

[18] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part i," *IEEE Robotics and Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006. 2, 3

[19] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part ii," *IEEE Robotics and Automation Magazine*, vol. 13, no. 3, pp. 108–117, 2006. 2, 3

[20] E. Ahokas, S. Kaasalainen, J. Hyypa, and J. Suomalainen, "Calibration of the optech altm 3100 laser scanner intensity data using brightness targets," *Int. Arch. Photogrammetry Remote Sensing Spat. Inf. Sci.*, vol. 34, pp. 3–6, 2006. 2

[21] B. Hofle and N. Pfeifer, "Radiometric calibration of terrestrial laser scanners with external reference targets," *Remote Sensing*, vol. 1, pp. 144–158, 2009. 2

[22] N. Pfeifer, P. Dorninger, A. Haring, and H. Fan, "Investigating terrestrial laser scanning intensity data: Quality and functional relations," in *Proceedings of the 8th Conference on Optical 3-D Measurement Techniques*, Zurich, Switzerland, July 9-12, 2007, pp. 328–337. 2

[23] N. Pfeifer, B. Hofle, C. Briese, M. Rutzinger, and A. Haring, "Analysis of the backscattered energy in terrestrial laser scanning data," *Int. Arch. Photogramm. Remote Sens.*, vol. 37, pp. 1045–1052, 2008. 2

[24] J. Levinson and S. Thrun, "Unsupervised calibration for multi-beam lasers," in *International Symposium on Experimental Robotics (ISER)*, 2010. 8, 9, 10, 14

[25] M. Unser and P.D. Tafti, "Stochastic models for sparse and piecewise smooth signals," *IEEE Transactions on Signal Processing*, vol. 59, no. 3, pp. 989–1006, March 2011. 3

[26] E. Bostan, U. Kamilov, M. Nilchian, and M. Unser, "Sparse stochastic processes and discretization of linear inverse problems," *IEEE Transactions on Image Processing*, vol. 22, no. 7, pp. 2699–2710, July 2013. 3

[27] A. Beck and M. Teboulle, "Fast gradient-based algorithm for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, November 2009. 6

[28] D.L. Donoho, "De-noising by soft thresholding," *IEEE Transactions on Information Theory*, vol. 41, pp. 613–627, 1995. 6, 11

[29] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1–4, pp. 259–268, November 1992. 6

[30] P. Perez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 313–318, July 2003. 6

[31] Y.E. Nesterov, "A method for solving convex programming problem with convergence rate $o(1/k^2)$," *Dokl. Akad. Nauk SSSR*, vol. 269, no. 3, pp. 543–547, 1983. 7, 11

[32] G. Pandey, J. R. McBride, and R. M. Eustice, "Ford campus vision and lidar data set," *IJRR International Journal of Robotics Research*, 2011. 7

[33] A.V. Segal, D. Haehnel, and S. Thrun, "Generalized-icp," *Robotics: Science and Systems*, 2009. 8

[34] M. Kaess, A. Ranganathan, and F. Dellaert, "isam: Incremental smoothing and mapping," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008. 8

[35] M. Kaess, H. Johannsson, R. Roberts, Viorela Lla, J. Leonard, and F. Dellaert, "isam2: Incremental smoothing and mapping using the bayes tree," *The international journal of Robotics research*, vol. 31, no. 2, pp. 216–235, 2012. 8

[36] J. Levinson and S. Thrun, "Automatic online calibration of cameras and lasers," in *Robotics: Science and Systems*, Berlin, Germany, June 24-28, 2013, pp. 29–36. 8

[37] A. Beck and M. Teboulle, "Fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, March 2009. 11

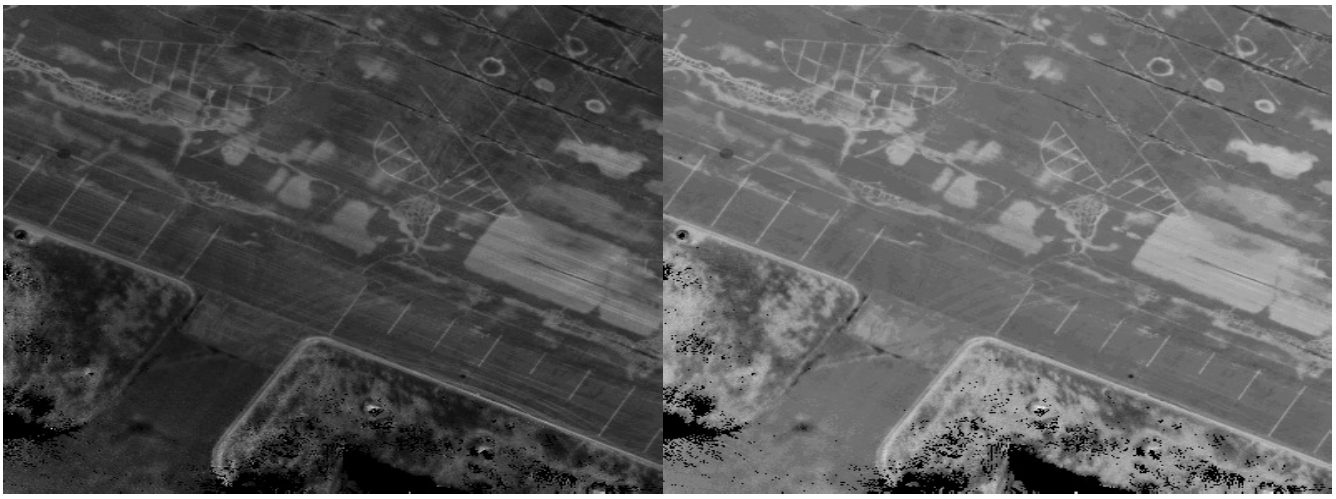

(a) Raw naïve fusion (b) Calibrated fusion

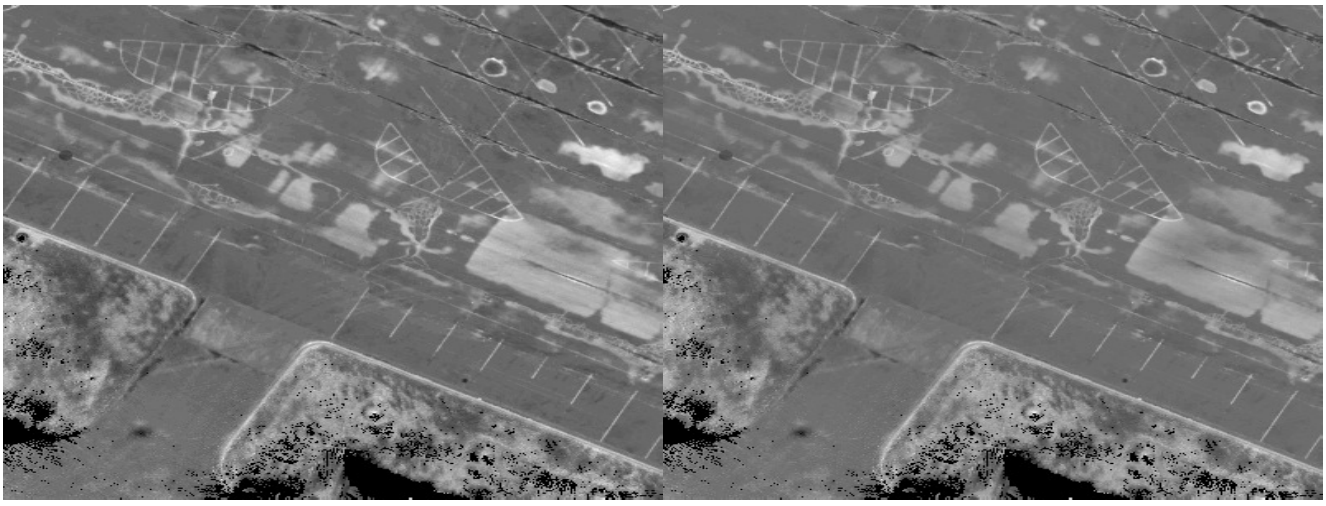

(c) Mean fusion no denoising (d) Mean fusion with denoising

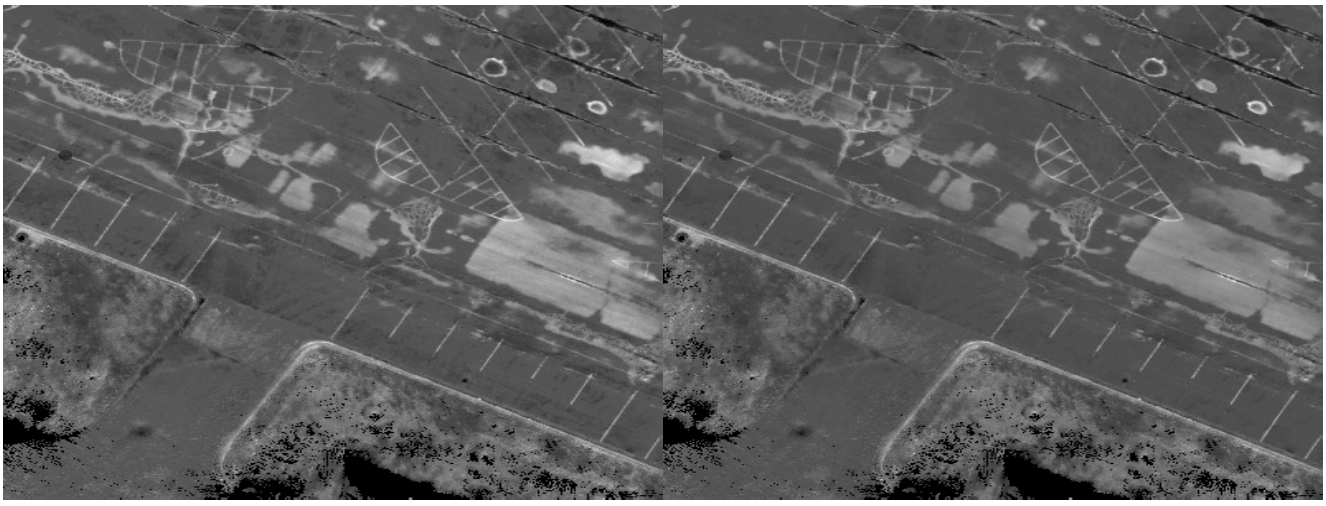

(e) Top 5 fusion no denoising (f) Top 5 fusion with denoising

Figure 7. Map patches of ground reflectivity. (a) Raw naïve fusion, (b) Our practical implementation of [24], (c) Our reconstruction using all map-perspectives, (d) same as (c) with denoising, (e) Our reconstruction with only five map-perspectives, (f) Same as (e) but with denoising. **Zoom in to better appreciate image quality.**